

## CLAIMS

What is claimed is:

- 1    1.    A method for mirroring data between a plurality of sites, comprising:
  - 2       establishing a replication relationship between the plurality of sites including a first site and a second site;
  - 4       replicating at least some changes made at any site of said plurality of sites at each other site of said plurality of sites;
  - 6       a first server associated with the first site requesting performance of a write operation;
  - 7       in response to said request, performing the write operation at the first site, and
  - 8       sending a message to request that the write operation be mirrored at the second site;
  - 10      the first server receiving an indication that the write operation could not be successfully performed at the second site;
  - 12      in response to the indication, the first server initiating a membership voting operation to determine which of the first server and a second server associated with the second site should be removed from said replication membership; and
  - 14      if it is determined that the second server associated with the second site should be removed from said replication membership, then the first server sending a message to indicate that write operations are no longer to be replicated at said second site.
- 1    2.    The method of Claim 1, further comprising the steps of:
  - 2       the second site rejoining the replication relationship; and

3       in response to said second site rejoining the replication relationship, causing data  
4               associated with the second site to be resynchronized with data that resides at  
5               the first site, and remounting a database associated with the second site at the  
6               second site.

1     3.   The method of Claim 1, further comprising the step of:  
2               if during the membership voting operation it is determined that the first site should be  
3               removed from said replication membership, then informing the second site  
4               that data is no longer to be replicated at the first site.

1     4.   The method of Claim 3, wherein the step of informing comprises:  
2               at the second site informing a file server associated with the second site or the second  
3               server that data is no longer to be replicated at the first site.

1     5.   The method of Claim 1, wherein the determination of which of the first server and the  
2               second server associated with the second site should be removed from said replication  
3               membership comprises determining which of the first server or the second server is  
4               more important or more reliable.

1     6.   The method of Claim 1, further comprising the step of establishing a particular device  
2               as a quorum device, and wherein the step of initiating a membership voting operation  
3               comprises the steps of:

4               notifying the quorum device that the write operation could not be successfully  
5               performed; and

6               the quorum device determining which of the first server and a second server  
7               associated with the second site should be removed from said replication  
8               membership.

- 1      7. The method of Claim 6, wherein the quorum device comprises a plurality of mirrored
- 2                devices, and the step of notifying the quorum device is performed by a primary file
- 3                server, wherein the primary file server is a file server, associated with one of the
- 4                plurality of sites, through which all other files servers associated with other sites in
- 5                the plurality of sites communicate with the quorum device.
  
- 1      8. The method of Claim 1, wherein the indication that the write operation could not be
- 2                successfully performed at the second site is an input/output error that is received at
- 3                the first server.
  
- 1      9. The method of Claim 1, further comprising the step of:  
2                if a particular site of the plurality of sites becomes inoperable, then initiating recovery  
3                of the particular site after it is determined that all messages transmitted from  
4                the particular site to each other site of the plurality of sites have been received  
5                at their destination.
  
- 1      10. A machine-readable medium carrying one or more sequences of instructions for  
2                mirroring data between a plurality of sites, wherein execution of the one or more  
3                sequences of instructions by one or more processors causes the one or more  
4                processors to perform the steps of:  
5                establishing a replication relationship between the plurality of sites including a first  
6                site and a second site;  
7                replicating at least some changes made at any site of said plurality of sites at each  
8                other site of said plurality of sites;  
9                a first server associated with the first site requesting performance of a write operation;

10           in response to said request, performing the write operation at the first site, and  
11                 sending a message to request that the write operation be mirrored at the  
12                 second site;  
13           the first server receiving an indication that the write operation could not be  
14                 successfully performed at the second site;  
15           in response to the indication, the first server initiating a membership voting operation  
16                 to determine which of the first server and a second server associated with the  
17                 second site should be removed from said replication membership; and  
18           if it is determined that the second server associated with the second site should be  
19                 removed from said replication membership, then the first server sending a  
20                 message to indicate that write operations are no longer to be replicated at said  
21                 second site.

1       11. The machine-readable medium of Claim 10, wherein execution of the one or more  
2           sequences of instructions by the one or more processors causes the one or more  
3           processors to further perform the steps of:  
4                 the second site rejoining the replication relationship; and  
5                 in response to said second site rejoining the replication relationship, causing data  
6                 associated with the second site to be resynchronized with data that resides at  
7                 the first site, and remounting a database associated with the second site at the  
8                 second site.

1       12. The machine-readable medium of Claim 10, wherein execution of the one or more  
2           sequences of instructions by the one or more processors causes the one or more  
3           processors to further perform the step of:

4 if during the membership voting operation it is determined that the first site should be  
5 removed from said replication membership, then informing the second site  
6 that data is no longer to be replicated at the first site.

1 13. The machine-readable medium of Claim 12, wherein the step of informing comprises:  
2 at the second site informing a file server associated with the second site or the second  
3 server that data is no longer to be replicated at the first site.

1 14. The machine-readable medium of Claim 10, wherein the determination of which of  
2 the first server and the second server associated with the second site should be  
3 removed from said replication membership comprises determining which of the first  
4 server or the second server is more important or more reliable.

1 15. The machine-readable medium of Claim 10, wherein execution of the one or more  
2 sequences of instructions by the one or more processors causes the one or more  
3 processors to further perform the step of:  
4 establishing a particular device as a quorum device, and  
5 wherein the step of initiating a membership voting operation comprises the steps of:  
6 notifying the quorum device that the write operation could not be successfully  
7 performed; and  
8 the quorum device determining which of the first server and a second server  
9 associated with the second site should be removed from said  
10 replication membership.

1 16. The machine-readable medium of Claim 15, wherein the quorum device comprises a  
2 plurality of mirrored devices, and the step of notifying the quorum device is  
3 performed by a primary file server, wherein the primary file server is a file server,

4       associated with one of the plurality of sites, through which all other files servers  
5       associated with other sites in the plurality of sites communicate with the quorum  
6       device.

1   17. The machine-readable medium of Claim 10, wherein the indication that the write  
2       operation could not be successfully performed at the second site is an input/output  
3       error that is received at the first server.

1   18. The machine-readable medium of Claim 10, wherein execution of the one or more  
2       sequences of instructions by the one or more processors causes the one or more  
3       processors to further perform the step of:  
4              if a particular site of the plurality of sites becomes inoperable, then initiating recovery  
5                  of the particular site after it is determined that all messages transmitted from  
6                  the particular site to each other site of the plurality of sites have been received  
7                  at their destination.